

Docket No. 007.0182.01
(OID 2000-167-01)

ORACLE CONFIDENTIAL
Patent

UNITED STATES PATENT APPLICATION

FOR

SYSTEM AND METHOD FOR PROVIDING A JAVA CODE RELEASE INFRASTRUCTURE WITH
GRANULAR CODE PATCHING

INVENTORS:
KENTON E. NOBLE
RICHARD J. LOTERO
SUBASH CHADALAVADA
WILLIAM GREENE
KEITH SWARTZ
JOSEPH H. KENT
JOSEPH EPSTEIN

PREPARED BY:

THE LAW OFFICES OF PATRICK J.S. INOUYE
810 THIRD AVENUE, SUITE 258
SEATTLE, WA 98104
(206) 381-3900

ASSIGNEE:
ORACLE CORPORATION
500 ORACLE PARKWAY
REDWOOD SHORES, CA

“Express Mail” mailing label number: EV029320346US

Date of Deposit FEB 11 2002

5 **SYSTEM AND METHOD FOR PROVIDING A JAVA CODE RELEASE
INFRASTRUCTURE WITH GRANULAR CODE PATCHING**

Cross-Reference to Related Application

This patent application claims priority under 35 U.S.C. § 119(e) to
provisional patent application Serial No. 60/268,886, filed February 14, 2001, the
10 disclosure of which is incorporated by reference.

Field of the Invention

The present invention relates in general to maintaining applications
developed using the Java programming language and, in particular, to a system
and method for providing a Java code release infrastructure with granular code
15 patching.

Background of the Invention

Over the last ten years, corporate computing environments have evolved
from monolithic proprietary mainframe computer systems to networks of
decentralized heterogeneous computer systems. These distributed systems
20 typically include personal computers and work stations that offer productivity
solutions with higher flexibility at lower cost.

At the same time, the cost and complexity of support has drastically
increased. For example, the operating systems and applications software require
periodic updates. Moreover, every different system configuration can require a
25 separate reinstallation of software, an operation which can be costly and time-
consuming for a large installed user base.

The Java programming language was developed to encourage low
maintenance application development by providing a machine-independent and
architecture-neutral programming language capable of operating on

heterogeneous computer systems independent of machine-specific environment and configuration settings. An overview of the Java programming language is described in P. van der Linden, "Just Java," Chapter 1, Sun Microsystems, Inc (1997 2d ed.), the disclosure of which is incorporated by reference.

5 The most common Java programs are applications and applets.

Applications are stand-alone programs. Applets are similar to applications, but adhere to a set of interface conventions and run only within the context of a Java-enabled container, such as a Web browser. During execution, the container downloads the byte code required to run the applet as needed.

10 The Java programming language provides a convenient means for disseminating software updates using Zip archives or Java Archive (JAR) files. These archives and files can be used to deliver and directly run Java applications and applets. The Zip archives and JAR files are packaged in a Zip format that allows multiple files to be bundled as a single archive file. In addition, JAR files
15 include a manifest storing attributes for each of the bundled files and can include an optional digital signature. JAR files are described in D. Flanagan, "Java in a Nutshell," Chapter 24, O'Reilly and Assocs. (1999 3d ed.), and the Zip format is described in <http://www.pkware.com/support/appnote.txt>, v. 4.5, PKWARE, Inc., Brown Deer, Wisconsin (May 1, 2001), the disclosures of which are incorporated
20 by reference.

Although offering a convenient way to deliver Java program updates, Zip archives and JAR files create a problem, known as the "high water mark" problem. The high water mark problem results from the inability to deliver specific updates or patches to individual component files using a JAR file.

25 Rather, when Java program updates or patches are delivered using JAR files, an entire JAR file is delivered, instead of just a single updated component file. Moreover, to ensure that each JAR file is current, the JAR file must always include the latest version of each component file. A single updated component file cannot be delivered and a complete high water mark JAR file must be
30 delivered instead. Delivering a complete high water mark JAR file results in delivering many other potentially undesirable changes, which can include

unexpected behavioral changes, or other highly destabilizing changes to the product.

To avoid the high water mark problem, Java program updates can be delivered in the form of individual files stored within an operating system 5 directory structure. Each separate update or patch can then be delivered to effect individual changes without unanticipated side effects. However, this approach poses difficulties with file naming conventions. The Java programming language supports long file and class names which are language features not supported on every operating system. Java source and class file names can contain special 10 characters, such as the dollar sign, which can cause portability concerns.

Similarly, the Java programming language uses the Unicode character set, an internationalized character set that may not be fully supported on all platforms, potentially causing further portability concerns.

These individual update or patch files also introduce directory nesting 15 concerns. Java packages are implemented as directories in a hierarchical file structure. Packages can be stored in deeply nested directories. However, the nesting depth can exceed the maximum permissible depth on a given platform. Consequently, these Java programming language-specific file naming and 20 directory storage conventions limit the operating systems that can directly support single file updating and patching.

JAR files lack integrity controls for ensuring the proper application of 25 updates and patches. The order in which updates and patches are applied should not ordinarily determine the end result of the application. However, Java updates and patches can be generated and applied in any order. Moreover, neither JAR files nor individual patch and update files include a revision control mechanism to ensure that out-of-date class files do not overwrite the latest versions.

Inner classes further complicate Java program updating and patching. A single Java source file that uses the Java “inner class” construct can produce 30 multiple Java class files. Inner classes make tracking out-of-date Java class files difficult.

Finally, other issues arise when patching and updating Java applets.

Applets are generally loaded into a Java-enabled browser in a JAR file format.

Some applets must be signed with a digital signature to work properly.

Consequently, the JAR file in which applets are delivered must be signed. If a

5 Java application being updated or patched requires a signature, either the signed high water mark JAR file must be delivered or the JAR file must be assembled and signed at the customer site. In addition to the problem inherent with delivering any high water mark file, delivering signed high water mark JAR files does not allow the integration of third-party code, that is, code not available when 10 the JAR files are created.

Therefore, there is a need for an approach for providing a Java code release infrastructure. Preferably, such an approach would resolve the high water mark, patch integrity and revision control problems while supporting long file names, special characters, extended character sets, deeply nested packages, and 15 inner classes.

There is a further need for an approach for providing granular patching of individual class files while maintaining the integrity of the source file to which the class belongs. Preferably, such an approach would include a facility to integrate third-party code into the patching process.

20

Summary of the Invention

The present invention provides a system and method for building and maintaining applications written in the Java programming language. Resource units are incorporated into a Java code release infrastructure. Each resource unit is created from a set of source or derived files, preferably written in the Java 25 programming language, and can include third-party code. The resource units include metadata that identifies and describes the contents. The version attribute of each resource unit can be derived from a revision string in the source file maintained by a revision control system. A set of Java code libraries, each including one or more resource units, is created and staged. Individual Java code 30 patches, each made up of one or more resource units, are generated from Java code patch definitions. A patch tool compares the resource units in the Java code

patches with those in Java code libraries and updates the resource units in Java code libraries as needed. If necessary, a set of Java archive (JAR) files, as defined by Java archive file definitions, are created or updated by extracting resource units from the Java code libraries. The JAR files are signed with a digital

5 signature if required. In addition, obsolete JAR files are deleted.

An embodiment of the present invention is a system and method for providing a Java code release infrastructure with granular code patching. One or more Java code patches including at least one resource unit are provided. The Java code patch contains only those resource units that need to be updated but can

10 optionally contain additional resource units. Each resource unit includes metadata and file components. The file components can include source Java files, derived Java files, and non-Java files. One or more Java code libraries, each including at least one resource unit, are patched. The metadata for each resource unit in the Java code patches is compared to the metadata for each such corresponding

15 resource unit in the Java code libraries. The file components in each resource unit in the Java code patches are merged into the Java code libraries for each such corresponding resource unit that is out-of-date. If needed, only those Java archive files with resource units that are out-of-date when compared to corresponding resource units in Java code libraries are created or updated, and signed with a

20 digital signature also if required.

A further embodiment is a system and method for patching staged code in a Java code release infrastructure. One or more staged Java code libraries are maintained in a staged code repository. Each Java code library includes at least one resource unit. Each resource unit includes metadata and file components.

25 The Java code patches are stored in a patch repository and can be downloaded to the client to stage. One or more Java code patches in a staged patch repository are accessed. Each Java code patch includes at least one resource unit. The metadata for each resource unit in the staged Java code patches is compared to the metadata in the staged Java code libraries for each such corresponding resource unit. The

30 file components in each resource unit in the staged Java code patches are merged into the staged Java code libraries for each such corresponding resource unit that

is out-of-date. The Java code libraries can be implemented as a virtual file system and can be used directly by a Java Virtual Machine (JVM). If necessary, a set of Java archive (JAR) files, as defined by Java archive file definitions, are created or updated by extracting resource units from the Java code libraries. The JAR files 5 are signed with a digital signature if required. In addition, obsolete JAR files are deleted.

A further embodiment is a system and method for generating Java code patches in a Java code release infrastructure. One or more Java code patch definitions describe the contents of Java code patches. A Java code patch 10 definition identifies the set of resource units that must be included in the corresponding Java code patch. Each resource unit includes metadata and file components and is generated using files maintained in a source code repository. The resource units are packaged into Java code patches and staged. The staged Java code patches are stored into a staged patch repository from which the 15 metadata for each resource unit is compared to the metadata for each such corresponding resource unit in the staged Java code libraries. The file components in each resource unit in the staged Java code patches are merged into the staged Java code libraries for each such corresponding resource unit that is out-of-date.

20 Still other embodiments of the present invention will become readily apparent to those skilled in the art from the following detailed description, wherein is described embodiments of the invention by way of illustrating the best mode contemplated for carrying out the invention. As will be realized, the invention is capable of other and different embodiments and its several details are 25 capable of modifications in various obvious respects, all without departing from the spirit and the scope of the present invention. Accordingly, the drawings and detailed description are to be regarded as illustrative in nature and not as restrictive.

Brief Description of the Drawings

FIGURE 1 is a functional block diagram of a distributed computing environment, including a system for providing a Java code release infrastructure with granular code patching, in accordance with the present invention.

5 FIGURE 2 is a block diagram showing the patch generator used in the system of FIGURE 1.

FIGURE 3 is a block diagram showing the patch tool used in the system of FIGURE 1.

10 FIGURES 4 and 5 are process flow diagrams showing prior art Java code patching systems.

FIGURE 6 is a process flow diagram showing the Java code release infrastructure system of FIGURE 2 and FIGURE 3.

15 FIGURE 7 is a block diagram showing the software modules of the resource unit generator of FIGURE 2.

FIGURE 8 is a tree diagram showing, by way of example, a sparse Java tree structure.

FIGURE 9 is a block diagram showing the software modules of the patch tool of FIGURE 3.

20 FIGURE 10 is a flow diagram showing a method for providing a Java code release infrastructure with granular code patching, in accordance with the present invention.

FIGURE 11 is a flow diagram showing the routine for building Java code patches for use in the method of FIGURE 10.

25 FIGURE 12 is a flow diagram showing the routine for creating resource units for use in the routine of FIGURE 11.

FIGURE 13 is a flow diagram showing the routine for creating Java code patches for use in the routine of FIGURE 11.

FIGURE 14 is a flow diagram showing the routine for patching staged code for use in the method of FIGURE 10.

30 FIGURE 15 is a flow diagram showing the routine for applying Java code patches for use in the routine of FIGURE 14.

Detailed Description

FIGURE 1 is a functional block diagram of a distributed computing environment, including a system 10 for providing a Java code release infrastructure with granular code patching, in accordance with the present invention. The system 10, as further described below beginning with reference to FIGURE 2, includes a patch generator 17 operating on a server 11 and a patch tool 21 operating on a client 12 or remote client 14. The server 11 is interconnected with the client 12, via a network 13, such as an intranetwork. The server 11 is coupled to a server storage device 18 while the client 12 is coupled to a client storage device 22. The server 11 can also be interconnected with a remote client 14 and coupled to a remote client storage device 23, via an internetwork 15, such as the Internet. The intranetwork 13 is interfaced to the internetwork 15 via a gateway 16. Other network topologies, configurations and components are feasible, as would be recognized by one skilled in the art.

The server 11 is used to build and maintain applications written in the Java programming language and other application programming languages. The applications are stored as source files 19 in the server storage device 18. As the applications are modified by updates and bug fixes, the patch generator 17 generates a set of Java code patches 20 that are downloaded by the client 12 and remote client 14. The patch tool 21 compares the contents of the Java code patch 20 to the contents of staged code 24 stored in the client storage device 22 or remote client storage device 23, respectively. The staged code 24 is updated by the patch tool 21 as needed.

The individual computer systems, including the server 11, client 12 and remote client 14, are general purpose, programmed digital computing devices consisting of a central processing unit (CPU), random access memory (RAM), non-volatile secondary storage, such as a hard drive or CD-ROM drive, network interfaces, and peripheral devices, including user-interfacing means, such as a keyboard and display. Program code, including software programs written in the Java programming language, and data, are loaded into the RAM for execution and

processing by the CPU and results are generated for display, output, transmittal, or storage.

FIGURE 2 is a block diagram 30 showing the patch generator 17 used in the system 10 of FIGURE 1. The patch generator 17 consists of two modules: 5 resource unit generator 31 and packager 32. The resource unit generator 31 processes a set of files 19 (hereafter “source”) to create resource units 38. The packager 32 bundles resource units 38 into a set of Java code patches 20, as defined by a set of Java code patch definitions 37. The Java code patch definitions 37 identify resource units 38 and the versions that must be included in the Java 10 code patches 20. The source 19 and the Java code patches 20 are stored in the server storage device 18. The Java code patches 20 are made available to the patch tool 21, as further described below with reference to FIGURE 3.

Individual Java applications (not shown) are executed by a JVM. Each Java application comprises byte code, which is a machine-independent form of 15 binary code and interpretable by any JVM independent of architecture or operating system. The byte code is included in the resource units 38.

Resource units 38 form the basis of the Java code release infrastructure to enable granular code patching. A resource unit 38 is the smallest entity capable of being updated or patched. Resource units 38 are formed from metadata 20 information (not shown) and file components stored in the source files 19, which include Java source code 33, Java derived code 34, non-Java source and derived code 35, and third party code 36. Metadata information is derived from the Java source code 33, Java derived code 34, non-Java source and derived code 35, and third party code 36, and includes attributes, such as a unique identifier, version 25 and similar information. A set of Java code patch definitions 37 are used by the resource unit generator 31 to identify the versions of the resource units 38 to include in each Java code patch 20.

FIGURE 3 is a block diagram 40 showing the patch tool 21 used in the system 10 of FIGURE 1. The Java code patches 20 are staged on client storage 30 device 22 or remote client storage device 23 (not shown). The patch tool 21 compares the set of Java code patches 20, packaged by the patch generator 17

(shown in FIGURE 2) to the set of staged code 24 and updates the staged code 24 as necessary. The resource units within the staged Java code libraries 41 are updated with resource units from the Java code patches 20, as further described below with reference to FIGURE 6. The set of staged code files 24 includes a set 5 of Java code libraries 41 and a set of Java archive files 42. The patch tool 21 maintains a set of Java archive file definitions 43 for describing the contents of the Java archive files 42 in terms of individual resource units 38.

Each Java code library 41 includes one or more resource units 38 bundled in a Java archive (JAR) or cabinet archive (CAB) format. A Java code library in 10 Zip format provides a portable virtual file system, which can be used directly by a Java virtual machine (JVM). The Java code libraries 41 can also be stored in an operating system directory structure or in a database (not shown).

The set of Java code libraries 41 need not be in the same format. Java code libraries 41 referenced by the patch tool 21 can be third-party code, which is 15 maintained separately from the Java code release infrastructure.

The Java archive files 42 can be used by stand-alone Java applications or Java applets running in a browser application. Each Java archive file 42 includes one or more resource units 38 plus additional functionality for use by Java applets. Each Java archive file 42 also includes a manifest storing attributes 20 describing the included resource units 38 and an optional digital signature.

Each Java archive file definition 43 describes the contents of an associated Java archive file 42 in terms of resource units 38. Similarly, the Java archive file definitions 43 can be used to describe an associated Java archive file that has become obsolete. The Java archive file definition 43 is extensible and allows 25 customizations.

Thus, when creating or updating Java archive files 42, the patch tool 21 first references the set of Java archive file definitions 43. The patch tool 21 compares the resource units 38 from the Java code libraries 41 with the corresponding resource units 38 from the Java archive files 42. Resource units 38 30 required in a Java archive file 42 can be from third-party Java code libraries 41 which are maintained separately from the infrastructure. The patch tool 21

updates the Java archive files 42 by replacing out-of-date resource units 38 from the Java archive files 42 with newer resource units 38 from the Java code libraries 41 as needed.

FIGURES 4 and 5 are process flow diagrams 50 and 70, respectively, 5 showing prior art Java code patching systems for patching JAR files and staged code in a directory structure, respectively. The respective process flows are divided into build phases 51, 71 and patch phases 52, 72.

Referring first to FIGURE 4, Java source code 33 is compiled (transition 54) using a Java compiler (not shown) into Java derived code 34. For example, a 10 Java source code file *A.java*, when compiled using a Java compiler, creates a Java class file *A.class*. The Java class file *A.class* is said to be derived from Java source code *A.java*. Similarly, a Java source file that uses an “inner class” structure compiles into multiple derived Java class files. Non-Java code includes, but is not limited to, files such as an GIF image file required by a Java application 15 or a Java applet. The Java source code 33 and Java derived code 34 along with non-Java source and derived code 35 and third-party code 36 required for Java applets are archived and compressed using the Zip format (transition 55) into a Java archive file 56. Optionally, the Java archive file 56 is signed (transition 57) with a digital certificate to form a signed Java archive file 58. The Java archive 20 file 56 or the signed JAR file 58 is packaged (transition 59) into a patch file 60. The patch file 60 is then staged (transition 61) as a staged Java archive file 62 and made available for patching (transition 52) for use by Java applications and applets (not shown).

Referring now to FIGURE 5, as before, Java source code 33 is compiled 25 (transition 54) using a Java compiler (not shown) into Java derived code 34. Those Java source files 33 and Java derived code 34 used by stand-alone Java applications are packaged (transition 59) into a patch file 60. The patch file 60 is then staged (transition 61) and made available for patching (patch phase 72) 30 directly into an operating system directory structure as staged code in the directory structure 62.

Referring back to FIGURE 4, compiling the Java source code files 33 and packaging the Java derived code 34 along with other non-Java code 35 and third-party code 36 into a single Java archive file 56 prior to staging causes the high water mark problem. Each Java archive file 56 must always be packaged with the 5 latest versions of the constituent Java source code files 33, Java derived code files 34, non-Java source and derived code files 35, and third-party code files 36. Otherwise, the Java archive file 56 could inadvertently back out changes previously made to Java source code files 33, Java derived code files 34, non-Java source and derived code files 35, or third-party code files 36 previously used to 10 patch staged code. The high water mark problem, causes any attempt to deliver a single change to potentially result in delivering other unanticipated or undesirable changes, including unexpected behavioral or destabilizing changes to an application or applet.

Similarly, when multiple Java class files 34 are derived from a single Java 15 source code file 33 such as required when updating Java inner classes, the prior art Java code patching system does not guarantee that all of the derived Java classes are updated as a unit. As well, the staged code in the directory structure 62 can place severe restrictions on the use of features supported by the Java programming language that may not be fully supported on all platforms. These 20 features include long file names, deeply nested packages, extended character sets, and special characters included in file names.

FIGURE 6 is a process flow diagram 90 showing the Java code release infrastructure system 30 of FIGURE 2 and system 40 of FIGURE 3. The process flow is divided into build phase 91 and patch phase 92. During the build phase 25 91, Java code patches 20 are generated along three individual processing threads (99, 100, 101) based on the type of originating source code. The set of Java code patch definitions 37 (shown in FIGURE 2) defines the contents of each Java code patch 20 in terms of resource units 38.

During the first thread 99, Java source code files 33 are compiled 30 (transition 54) using a Java compiler (not shown) into Java derived code files 34. Resource units 38 are generated (transition 95) from the source 33 and Java

derived code files 34, as further described below with reference to FIGURE 7. This set of files 33 and 34 can include Java class files that correspond to Java inner classes and use Java-native file naming conventions. Metadata information (not shown) for the source and derived Java code files 33,34 is generated

5 (transition 95) and resource units 38 are packaged (transition 96) into Java code patches 20. The metadata information for each resource unit 38 identifies the contents of the resource unit 38 and includes attributes, including a unique identifier and version number. The resource unit version can be derived from a revision string in each Java source code file 33 maintained by a revision control

10 system (not shown).

During the second thread 100, non-Java source and derived code files 35 are used to generate (transition 95) resource units 38, as further described below with reference to FIGURE 7. Metadata information (transition 95) is generated and the resource units 38 are packaged (transition 96) into the Java code patches

15 20.

During the third thread 101, third party code files 36 are used to generate (transition 95) resource units 38, as further described below with reference to FIGURE 7. Metadata information is generated (transition 95) and the resource units 38 are packaged (transition 96) into the Java code patches 20.

20 During the patch phase 92, the Java code patches 20 are compared to and merged into (transition 97) Java code libraries 41 using the system of FIGURE 3. Resource units in the Java code patches 20 are compared to the corresponding resource units in the Java code libraries 41 and, if necessary, merged (transition 97). To merge the resource units, the out-of-date resource units in the Java code

25 libraries 41 are replaced by the corresponding resource units in the Java code patches 20. In addition, any resource units missing from the Java code libraries 41 are added from the Java code patch 20 and any obsolete resource units are optionally deleted. Then, if necessary, resource units are extracted (transition 98) from the Java code libraries 41 and are used to create or update the Java archive

30 files 42. Finally, the Java archive files 42 are optionally signed (transition 98) using a digital certificate (not shown) to create signed Java archive files 42.

FIGURE 7 is a block diagram 120 showing the software modules of resource unit generator 31. The resource unit generator 31 is part of the patch generator 17 of the Java code release infrastructure system 10 of FIGURE 1. As described above with reference to FIGURE 2, the patch generator 17 includes two modules, resource unit generator 31 and packager 32. The patch generator 17 generates Java code patches 20 (shown in FIGURE 2) containing one or more resource units 38. Each resource unit 38 consists of a set of Java files, including Java source 33 and derived files 34, non-Java source and derived files 35, third-party code files 36, and metadata 124 describing the contents of the resource unit 38. The non-Java source files 35 can include, for instance, image files.

During operation, the resource unit generator module 31 identifies the set of files required for each resource unit 38 using the Java code patch definitions 37 (shown in FIGURE 2). Java source files 33 are compiled using a Java compiler 121 to produce a set of Java derived code 34. A single Java source file 33 can result in more than one Java derived file 34.

For example, the following Java code segment defines a Java class *A.class* and an inner Java class *A\$Inner.class*:

```
20      package oracle.apps.wip.util ;
          import java.io.* ;

          public class A {
              public static final String RCS_ID =
                  "$Header: A.java 115.3 2000/02/14 12:00:00 \
          25                  appldev ship $" ;

              private boolean done ;

              public A(final InputStream istream)
          30              {
                  Thread reader = new Thread ( ) {
                      public void run ( )
                      done = true ;
                      return A.foo() ;
          35                  }
              } ;
          } ;
          private static class Inner { int x ; int y ; }
      } 
```

40

When compiled, this Java source would result in multiple Java derived files 34, namely, *A.class*, *A\$1.class* and *A\$Inner.class*. These files must be packaged and updated as a unit.

FIGURE 8 is a tree diagram showing, by way of example, a sparse Java tree structure 130. The sparse Java tree 130 illustrates the hierarchical structuring of the Java code segment described above. The sparse Java tree 130 is a hierarchical data structure with a root node *JAVA_TOP* 131. The individual Java classes *A.class*, *A\$1.class*, and *A\$Inner.class*, are defined in a Java source code file and correspond to components of a resource unit staged together to reflect class inheritance. The Java class *A.class* 132 has two child classes, *A\$1.class* 133 and *A\$Inner.class* 134. When compiled, this Java class definition would result in multiple class files, namely, *A.class* 132, *A\$1.class* 133, and *A\$Inner.class* 134.

Referring back to FIGURE 7, after all Java source files 33 have been compiled, the metadata generator 122 generates metadata information 124 that is included in the resource units 38. The metadata information 124 describes the contents of each resource unit 38 and includes other attributes, such as an identifier that uniquely identifies the resource unit 38 and a version attribute used when comparing resource units.

The resource unit packager 123 bundles the Java source 33, Java derived code 34, non-Java source and derived code 35, third-party code 36, and metadata 124 into the resource units 38. In the described embodiment, resource units 38 are packaged as Zip archives, although other formats are feasible, as described above. Resource units 38 allow the use of Java-native features, including long file and class names, file names containing special characters, Java inner classes, and deeply nested packages implemented as directories.

FIGURE 9 is a block diagram 140 showing the software modules of the patch tool 21 of FIGURE 3. The patch tool 21 updates the staged code files 24 (shown in FIGURE 1) using the resource units stored in Java code patches 20. The staged code 24 includes Java code libraries 41 and Java archive files 42. Each Java code library 41 can include Java source code, Java derived code, non-Java code and third-party code. Although the patch tool 21 is used to maintain

Java code libraries 41, not all Java code libraries need to be maintained by the Java release infrastructure. This flexibility allows for staged third party code. When generating the Java archive files 42, not necessarily all of the Java code libraries are referenced by the patch tool 21.

5 The patch tool 21 includes four modules: compare 141, merge 142, extract 143, and sign 144. The compare module 141 compares resource units included in the set of Java code patches 20 to corresponding resource units included in the Java code libraries 41. The compare module 141 compares, for instance, the version attribute of corresponding resource units to determine if a resource unit in 10 any Java code library 41 requires updating. The compare module uses a set of rules that allow a newer resource unit to be replaced by an older resource unit for the purposes of backing out a previously applied Java code patch.

15 The merge module 142 updates the Java code libraries 41 by removing obsolete resource units, adding any resource units that are missing and replacing out-of-date resource units in the Java code libraries 41 with updated resource units from the Java code patches 20.

20 In the described embodiment, any resource unit included in a Java code library 41 that does not have a version attribute is replaced with a corresponding resource unit from a Java code patch 20. The set of rules used when comparing resource units and determining out-of-date resource units are flexible and provide the ability to back out changes made by a prior Java code patch, when required. Other merging rules can be applied to implement a resource unit updating policy, as would be recognized by one skilled in the art.

25 The extract module 143 uses the set of Java archive file definitions 43 to identify a set of Java archive files 42 in terms of included resource units. Each Java archive file 42 includes a set of resource units stored in a Java archive (JAR) file format. Java archive files 42 are primarily used by Java applets running in Java-enabled browser applications. If a required Java archive file 42 does not already exist as a staged code file 24 or requires an update, the extract module 143 30 extracts the required resource units from the Java code libraries 41 to create or update the Java archive file 42.

Finally, if required the sign module 144 digitally signs Java archive files 42 using a digital certificate 145. Digital signatures are used by Java applications and applets to authenticate the sources of Java code included in Java archive libraries.

5 FIGURE 10 is a flow diagram showing a method 150 for providing a Java code release infrastructure with granular code patching, in accordance with the present invention. The method proceeds in two phases: build 151 and patch 152, as further described below in FIGURE 11 and FIGURE 14, respectively. During the build phase 151, Java code patches 20 (shown in FIGURE 2) are created by
10 the patch generator 17 (shown in FIGURE 2). During the patch phase 152, the resource units 38 (shown in FIGURE 3) stored in the Java code patches 20 are compared and, if necessary, merged into the contents of the staged code 24 using the patch tool 21 (shown in FIGURE 3).

15 FIGURE 11 is a flow diagram showing the routine 160 for building Java code patches 20 for use in the method 150 of FIGURE 10. The purpose of this routine is to create resource units which are packaged in Java code patches 20.

20 In the initialize phase (block 161), Java code patches 20 (shown in FIGURE 2) are identified for processing by checking the set of Java code patch definitions 37. As necessary, resource units 38 are created (block 162), as further described below with reference to FIGURE 12. Finally, the Java code patches 20 are created (block 163), as further described below with reference to FIGURE 13. The routine then returns.

25 FIGURE 12 is a flow diagram showing a method 170 for creating resource units 38 (shown in FIGURE 2) for use in the routine 160 of FIGURE 11. The purpose of this routine is to generate a set of resource units 38 from a set of source files 19 (shown in FIGURE 1).

30 The set of resource units 38 to be generated is identified by referring to the Java code patch definitions 37 (block 171). Each identified resource unit 38 is processed in an iterative loop (blocks 172-178). During each iterative loop (block 172), the resource unit 38 definition is first checked for third-party code (block 173) and native Java source code (block 174). If the resource unit does not

contain third-party code (block 173) but does contain native Java source code (block 174), the Java source code is compiled (block 175) using a Java compiler 121 (shown in FIGURE 7). Metadata information 124 is generated by the metadata generator 122 (block 176). Finally, the resource unit 38 is packaged 5 (block 177) using resource unit packager 123 (shown in FIGURE 7). Processing continues with the next resource unit (block 178), until all resource units have been processed (blocks 172-178). The routine then returns.

FIGURE 13 is a flow diagram showing a routine 180 for creating Java code patches 20 (shown in FIGURE 2) for use in the routine 160 of FIGURE 11.

10 The purpose of this routine is to create Java code patches 20 as defined by the set of Java code patch definitions 37.

The set of Java code patch definitions 37 is referenced (block 181). Each resource unit 38 referenced in the Java code patch definitions 37 is processed in an iterative loop (blocks 182-184). During each iterative loop (block 182), 15 resource units 38 are packaged into a Java code patch 20 (block 183) using the packager 32 of the patch generator 17 (shown in FIGURE 2). The iterative processing continues with each Java code patch 20 (block 184) until all of the Java code patches have been packaged (blocks 182-184). The routine then returns.

20 FIGURE 14 is a flow diagram showing a routine 190 for patching staged code for use in the method 150 of FIGURE 10. The purpose of this routine is to apply Java code patches 20 (shown in FIGURE 3) to staged code 24 (shown in FIGURE 3).

Initially, the Java code patches 20 are identified for processing (block 25 191). The Java code patches 20 are then applied (block 192), as further described below with reference to FIGURE 15. The routine then returns.

FIGURE 15 is a flow diagram showing a routine 200 for applying Java code patches 20 (shown in FIGURE 3) for use in the routine 190 of FIGURE 14. The purpose of this routine is to merge the resource units 38 included in a set of 30 Java code patches 20 into a set of staged Java code libraries 41 and, as required, Java archive files 42 (shown in FIGURE 3).

The routine merges a set of Java code patches 20 using a set of nested iterative processing loops (blocks 201-207 and 202-206). During each outer iterative loop (blocks 201-207), a Java code patch 20 is processed. During each inner iterative loop (blocks 202-206), a resource unit 38 included in the Java code patch 20 is processed. Resource units 38 in a Java code patch 20 are compared to corresponding resource units in Java code libraries 41 (block 203). If a Java code library 41 requires updating (block 204), the resource unit 38 in the Java code library 41 is updated (block 205). Processing continues with each subsequent resource unit 38 in the Java code patch 20 (block 206) and for each remaining Java code patch 20 (block 207). The routine then returns if Java archive files 42 are not required (block 208).

Otherwise, if Java archive files are required (block 208), the set of Java archive file definitions 43 (shown in FIGURE 3) are iteratively processed (blocks 209-216). During each iterative loop (block 209), the metadata information in Java archive files 42 and Java code libraries 41 is read (blocks 210 and 211, respectively). If the Java archive file 42 requires updating (block 212), the Java archive file 42 is updated (or created, if necessary) by extracting resource units 38 from the appropriate Java code libraries 41 (block 213). Updating is necessary when any resource unit 38 in the Java archive file 42 is out-of-date. Also, if a Java archive file 42 is obsolete (block 212), the Java archive file 42 is deleted (block 213). If the Java archive file 42 requires authentication by digital signature (block 214), the Java archive file 42 is digitally signed using a digital certificate 145 (shown in FIGURE 9) (block 215). Processing continues with each Java archive file definition 42 (block 216) until all Java archive file definitions 42 have been processed (blocks 209-216), after which the routine returns.

While the invention has been particularly shown and described as referenced to the embodiments thereof, those skilled in the art will understand that the foregoing and other changes in form and detail may be made therein without departing from the spirit and scope of the invention.